

REMARKS/ARGUMENTS

Minor changes are made to this specification. Claim 22 is amended. Claims 1-65 are pending in the application. Reexamination and reconsideration of the application, as amended, are respectfully requested.

The specification is amended to reflect the current status of certain related applications.

Claims 1, 4-10, 16-21, 24, 27-33, 39-44, 45, 48-54 and 60-65 were rejected as being anticipated by Smith (US 6347359). Claims 2, 3, 22, 23, 25, 26, 46, and 47 were rejected as being obvious over Smith taken in view of Uysal (US 6922752). Claims 11, 12, 34, 35, 55 and 56 were rejected as being obvious over Smith taken in view of Anderson (US 2002/0059539). Claims 13, 36 and 57 were rejected as being obvious over Smith taken in view of Belhadj (US 6516425). Claims 14, 15, 37, 38, 58 and 59 were rejected as being obvious over Smith taken in view of Schultz (US 6058489). These rejections are respectfully traversed.

The present invention is directed to a RAID conversion method. One feature of the invention claimed in independent claims 1, 24 and 45 is a rule-based method of determining the target RAID (i.e. the RAID that the disk array is to be converted to). Claim 1 recites: “storing in the controller one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array; detecting the one or more conditions of the array; selecting the desired RAID algorithm based on the detected conditions and the stored rules; ...” (Emphasis added.) Claim 24 recites: “the controller further storing one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array, the controller having stored program instructions or a logic circuit operable to detect the one or more conditions of the array, to select the desired RAID algorithm based on the detected conditions and the stored rules, ...” (Emphasis added.) Claim 45 recites: “first computer readable program code configured to cause the controller to storing one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array; second computer readable program code configured to cause the controller to detect the one or more conditions of the array; third computer readable program code configured to cause the controller to select the desired RAID algorithm based on the detected conditions and the stored rules; ...” (Emphasis added.)

Indeed, converting from one particular RAID to another particular RAID was known, as the applicant discussed in the “Background” section of the instant application. The inventive

features of the invention of claims 1, 24 and 45 lies in the use of rules to select the target RAID. Some examples of the rules are described in the specification at page 9, line 1 to page 11, line 11, including capacity utilization-based rules, performance requirement-based rules, self-healing fixed insurance, self-healing minimal insurance, data criticality-based rules, data recency and repetition-based rules, device vulnerability-based rules, etc. Of course, the invention is not limited to any particular rule being used.

None of the cited references teach or suggest a rule-based selection of the target RAID. The primary reference, Smith, merely discloses an implementation of a RAID conversion method. The flow charts in Figs. 2-6 show how the conversion should be carried out for various source/target RAID combinations, but the reference does not touch upon how the target RAID is selected, i.e., how it is determined what the target RAID should be. The Examiner relied on col. 6, lines 8-16 for teaching selecting the desired RAID algorithm based stored rules, states: "In the case where array controller 106 determines that source and destination configurations are not both RAID level zero, one disk configurations, execution of method 300 continues at step 310. In step 310, array controller 106 determines if the source configuration is RAID level zero, one disk and the destination configuration is RAID level one, two disk. Array controller 106 also determines if the single disk 104 included in the source configuration is included in the destination configuration." Clearly, in this step, the destination RAID has already been selected, and the array controller 106 performs different tasks depending on what the selected source and destination configurations are. Thus, Smith does not teach or suggest the above-quoted elements of claims 1, 24 and 45.

The secondary references do not cure the deficiency of the primary reference. Uysal describes a "computer storage system [that] includes a controller and a storage device array. The storage device array may include a first sub-array and a fast storage device sub-array. The first sub-array includes one or more first storage devices storing data. The fast storage device sub-array includes one or more fast storage devices storing a copy of the data stored in the first sub-array." (Abstract.) It is not related to RAID conversion. Anderson describes "a hybrid data reconstruction system and method for a data storage device. Data is selectively stored according to one of two or more redundancy schemes such that critical data is stored according to a scheme which has a higher degree of redundancy." (Abstract.) It is not related to RAID conversion either. Belhadj describes a "method of managing data in a hierarchical data storage system

employing data redundancy schemes includes prioritizing a data rebuild based on a most vulnerable data redundancy scheme in the storage system. ... Prioritizing the data rebuild includes enabling a rebuild of the most vulnerable data redundancy scheme prior to enabling a rebuild of any other data redundancy scheme in the system. The most vulnerable data redundancy scheme is determined by comparing a probability of losing data that can be prevented by a rebuild for each data redundancy scheme with respect to the potential for one or more next storage device failures in the data storage system. ..." (Abstract.) This reference relates to a data rebuilt but not to RAID conversion. As they do not relate to RAID conversion, the Uysal, Anderson and Belhadj references do not teach or suggest the above-quoted elements of claims 1, 24 and 45.

Schultz describes an on-line disk array reconfiguration method. It describe that the host issues a SET CONFIGURATION or a CHANGE CONFIGURATION command to reconfigure the disk array. These commands may include parameters that specify the mapping scheme used between a logical request and a physical request, e.g., RAID 0, RAID 1, RAID 4, RAID 5, etc. (col. 10, lines 13-47; see also col. 23, lines 10-36). It also describes setting a fault tolerance mode of the logical volume, e.g. RAID 4, RAID 1 and RAID 5 (col. 10, lines 54-60). The reference does not disclose, however, how the mapping scheme or the fault tolerance values (i.e., RAID 0, RAID 1, RAID 4, etc.) are selected, or whether the selection is based on any rule or logic. The reference merely indicates that these parameters are set by the user (see col. 11, lines 11-22). Thus, Schultz does not teach or suggest the above-quoted elements of claims 1, 24 and 45.

In short, none of the cited references teaches or suggests the above-quoted elements of claims 1, 24 and 45. Accordingly, these claims, as wee as their depend claims 2-21, 25-44 and 46-65 are patentable over the cited references.

Regarding independent claim 22, this claim has been amended to more specifically recite a RAID conversion method. The amended claim 22 is supported by the instant specification at, for example, page 8, lines 1-21.

The applicant believes that amended claim 22 is patentable over the cited references. For example, the Uysal reference, which is relied on by the examiner for teaching the use of a watermark, only mentions the use of watermarks in a general statement (e.g. col. 6, lines 1-2), but fails to teach or suggest many elements of the amended claim 22, including: "alternating

between performing tasks in a conversion state and performing tasks in a user I/O state” and “wherein the tasks in the user I/O state include reading data from the storage devices upon user requests, wherein when requested data exist on both sides of the watermark, the requested data is read according to the second RAID algorithm before the watermark and according to the first RAID algorithm after the watermark.” Further, Uysal’s approach would give very different results if used in RAID conversion. Uysal states (col. 6, lines 1-2): “A (potentially fault-tolerant) bitmap or watermark pointers [sic] can be used to keep track of the progress of the reconstruction. As in the case of copying to restore redundancy, this copying can occur through a memory buffer, allowing asynchronous reading and writing in the background.” Allowing asynchronous reading and writing merely on the basis of a memory buffer is vulnerable to “write hole” and other race condition problems, where old data supersedes new at unpredictable times. The method of the instant claim 22, on the other hand, can avoid write holes by alternating states.

The Schultz reference (not relied on by the Examiner for teaching a watermark) describes a concept of a boundary between already reconfigure data and not yet reconfigured data, although the term “watermark” is not used. In Schultz, however, the user read requests during reconfiguration is handled differently from that of the instant claim 22. Schultz describes: “While the reconfigure process is in progress, read and write requests from the host computer system are handled by the disk controller, which performs a check to determine if the request is to blocks that have already been reconfigured. If so, the read or write request is performed using the destination logical volume configuration information. If the accessed blocks have not yet been reconfigured, the source logical volume configuration information is used to perform the request. A request may also span blocks that have not yet been reconfigured and those that have been reconfigured. In those instances, the host request is suspended until the disk array reconfiguration process is able to catch up.” (Col. 3, lines 15-27, emphasis added.) Thus, it is clear that Schultz cannot handle request for data that exist across the boundary between the already reconfigured data and not yet reconfigured data. Claim 22, on the other hand, recites: “wherein when requested data exist on both sides of the watermark, the requested data is read according to the second RAID algorithm before the watermark and according to the first RAID algorithm after the watermark.”

Accordingly, the applicant believed that amended claim 22 and its dependent claim 23 are patentable over the cited references.

The applicant concurrently submits an Information Disclosure Statement with the requisite fee under 37 C.F.R. 1.97(c). The references cited in the IDS were discussed in the application itself in the "Background" section of as related patents/applications.

In view of the foregoing, it is respectfully submitted that the application is in condition for allowance. Reexamination and reconsideration of the application, as amended, are requested. If for any reason the Examiner finds the application other than in condition for allowance, the Examiner is invited to call the undersigned attorney at the Los Angeles, California telephone number (213) 625-5076 to discuss the steps necessary for placing the application in condition for allowance.

If there are any fees due in connection with the filing of this response or deficient in fees, please charge the fees to our Deposit Account No. 50-3531.

Respectfully submitted,

Date: May 23, 2007

By: /Ying Chen/
Ying Chen
Registration No. 50,193
Attorney for Applicant(s)

255 S. Grand Ave., #215
Los Angeles, CA 90012
Phone: 213-625-5076
Fax: 213-625-0691